Chapter 10

Front End Electronics

10.1 Introduction

The data acquisition systems of high rate HEP experiments typically rely on a "Level-1" trigger to reject most data before it is read out of the front-end electronics. In these experiments, data from a small fraction of the detector are read out quickly and input to the trigger system. After a relatively short and fixed length of time (typically less than a few microseconds), a trigger decision is made. During this time, data are stored in the front-end electronics or on passive elements such as delay cables. Only after a Level-1 accept occurs is most of the data read out of the front-end electronics.

By contrast, the BTeV Level-1 trigger system will take a relatively long time to make its decisions. Moreover, the time required by the Level-1 trigger will vary significantly from crossing to crossing; most crossings will be processed by the trigger system in hundreds of microseconds, but some may take orders of magnitude longer. This long, and variable, trigger latency makes it impractical to store data in the BTeV front-end electronics. Instead, all data from the entire BTeV detector, for every beam crossing, will be digitized, zero suppressed, and read out into buffer memory.

The front-end electronics associated with the different elements of the BTeV spectrometer share an architecture (see Figure 10.1). Data are digitized and zero-suppressed in electronics mounted on the detector and/or in electronics located in racks very close to the detector. Data from a number of front-end elements are collected by modules called Data Combiner Boards (DCB's) that are also located in the collision hall. The DCB's transmit the data over optical fiber links to the counting house. Some of the data streams are input directly into Level-1 Buffers; other data streams undergo another step of reformatting before being input to the Level-1 Buffers. Two types of custom serial links are used to transport data from on-detector or near-detector electronics and the DCB's. One type of link is used by the pixel and silicon strip detectors. The other type of link is used by the remaining subsystems.

The Data Combiner Boards are located in 6U Eurocard subracks. A "Clock Distributor" module, also located in the collision hall, distributes a 7.6 MHz (1/132 ns) clock over equal time copper links to the DCB's. In normal operation, this clock is derived from the

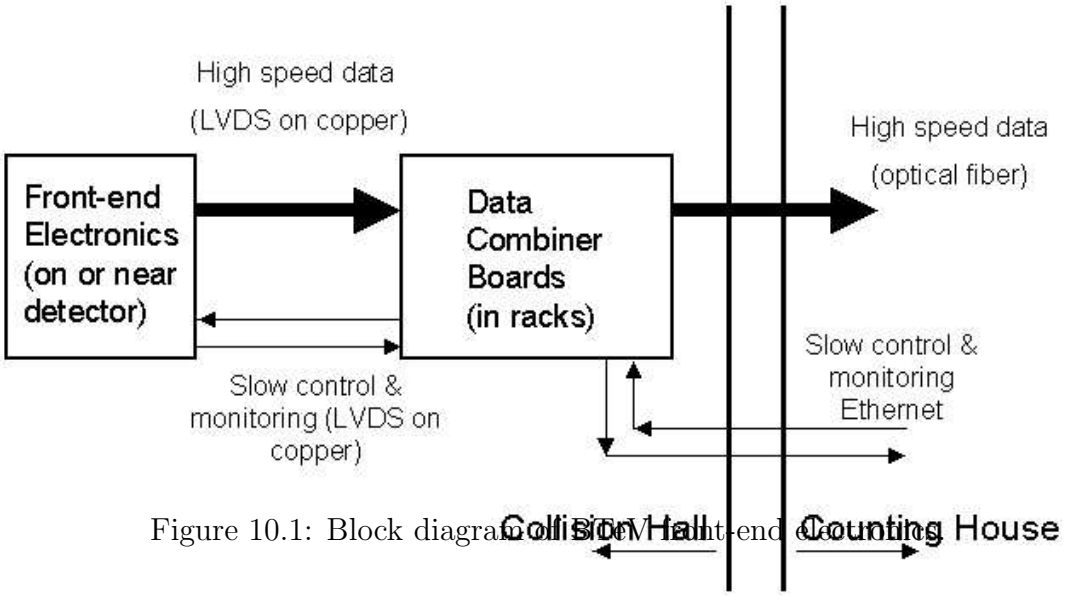

Figure 10.1: Block diagram of BTeV pixel front-end electronics.

Accelerator clock. Equal time point-to-point links between the Clock Distributor and the DCB's are also used for a "synchronize" signal that is used to synchronize various commands given to the DCB's. The Clock Distributor and the DCB's use Ethernet for slow control and monitoring communications.

10.2 Pixel Detector

10.2.1 Overview

The basic building block of the pixel detector is the module. A module consists of a silicon sensor bump bonded, depending on its size, to 4, 5, 6, or 8 FPIX2 readout chips. The chips are mounted on a High Density Interconnect (HDI) flexible printed circuit, and wire bonded to it. FPIX2 readout chips communicate with pixel DCB's using Low Voltage Differential Signaling (LVDS) over copper serial links. All of the FPIX2 chips attached to an HDI share one slow control and monitoring link, as well as common digital and analog supply voltages and grounds. Hit data is output from the FPIX2 chips on 140 Mbps point-to-point links. Chips nearest the beam are configured to use six data output links. Those further away from the beam are configured to use 4, 2, or 1 data output link.

Signals are carried on lightweight flexible cables between the HDI's and printed circuit boards on the sides of the pixel vacuum box. These printed circuit "feed through" boards carry the signals across the vacuum seal. Signals are carried between the feed through boards and the DCB's on conventional high-density cables ($\sim$10 m long). The DCB's are located in racks mounted on the outside of the return yoke of the BTeV dipole magnet. Two DCB's are used for each half-station.

10.2.2 Component Quantities and Locations

Table 10.1 summarizes the number of sensor modules, readout chips, DCB's, DCB subracks, and data links used in the pixel detector. Each of the sixty half-stations in the pixel detector contains one "x" half-plane and one "y" half-plane. "X" half-planes have the pixels oriented so that a precision position measurement is made in the x (non-bend) direction; "y" half-planes have the pixels oriented to make precision y measurements. A block diagram of the system is given in Figure 10.2.

	Number per "x" half-plane	Number per "y" half-plane	Number in complete detector
4-chip sensor modules	2	0	120
5-chip sensor modules	0	9	540
6-chip sensor modules	0	7	420
8-chip sensor modules	5	0	300
FPIX2 IC's	48	87	8100
DCB's	see text	see text	120
DCB subracks	NA	NA	10
140 Mbps data links to DCB's	108	149	15420
2.5 Gbps data links from DCB's	NA	NA	960

Table 10.1: Component Count

10.2.3 Data Structure

Each pixel hit output from an FPIX2 chip consists of 23 bits of data: an 8-bit beam crossing number, a 7-bit row number, a 5-bit column number, and a 3-bit pulse height. One additional bit is used to mark word boundaries, so 24 bits are transmitted per pixel hit. The DCB's add a 7-bit chip ID number to each hit and extend the beam crossing number to 11 bits before sending the data on commercial optical fiber links upstairs to the counting house. The data output by the FPIX2 chips is not strictly time-ordered. Time order is restored, the crossing number is further extended, and pixel data is reformatted into 16-bit words by Pixel Preprocessor boards located in the counting house. The reformatted data is stored in Level-1 Buffers and input to the Pixel Trigger Processor.

10.2.4 Occupancy and Data Rate Estimate

Detailed simulations have been run to verify that the FPIX2 readout chips will have adequate output bandwidth, even at luminosities much higher than the BTeV design luminosity of 2×10^{32} cm$^{-2}sec^{-1}$. These simulations used events generated by PYTHIA and BTeV GEANT. They included charge sharing in the pixel sensors due to geometry (track angle with respect to the detector) and charge carrier diffusion. They also included photon conversions and

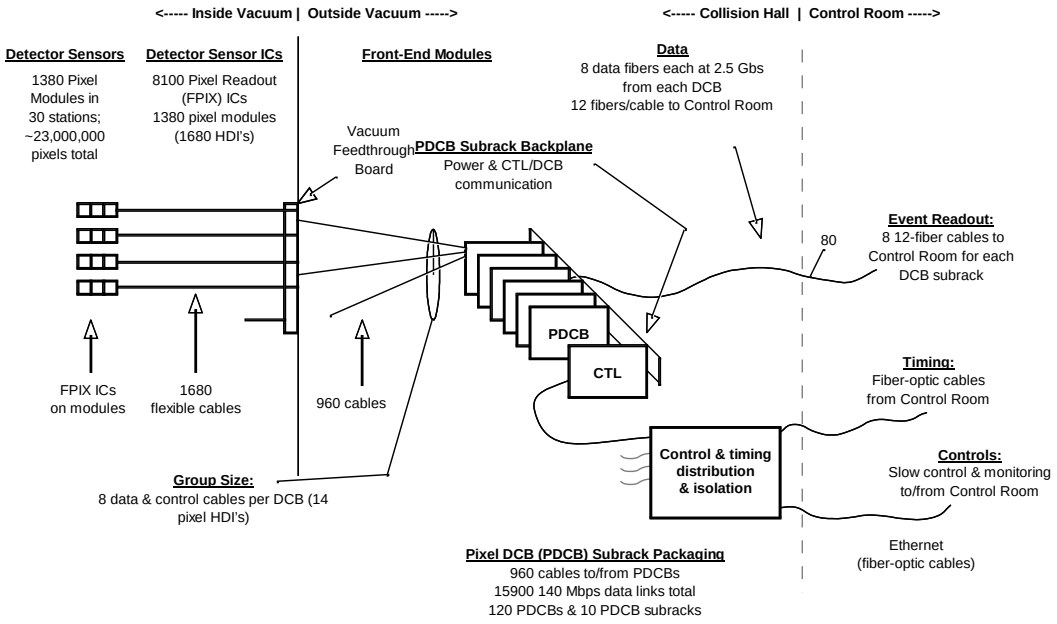

Figure 10.2: Block diagram of BTeV Pixel Detector front-end electronics.

interactions in material. Low energy delta rays (below the threshold used in GEANT) were not included in the simulations, but our 1999 test beam data indicate that this is a small effect.

The simulations show that, on average, each interaction generates just under 0.6 pixel hits in the 128 row by 22 column region of the central station closest to the beam that is covered by a single FPIX2 readout chip. The same simulations show that the occupancy falls off like $r^{-1.6}$. The average occupancy is approximately 0.1 pixel hits per interaction per FPIX2 readout chip.

Regardless of the operating mode of the Tevatron, the BTeV design luminosity of 2×10^{32} $cm^{-2}sec^{-1}$ corresponds to a peak interaction rate of approximately 15.2 MHz. Therefore, the data rate into the pixel DCB's is estimated to be:

$$0.1(\frac{hits/interaction}{FPIX2}) \times 8100(FPIX2's) \times 24(\frac{bits}{hit}) \times 15.2 \times 10^6(\frac{interactions}{sec}) = 0.3Tbps. \quad (10.1)$$

This is to be compared with the aggregate output bandwidth of the FPIX2 readout chips of more than 2 Tbps. At design luminosity, the data rate into each DCB is approximately 2.5 Gbps. The data rate from the DCB's to the pixel preprocessors is $\sim 40\%$ higher, since 10 bits are added to each 23 bit data word. Data will be transmitted on commercial optical fiber links from the DCB's to the pixel preprocessors. The final choice of link has not yet been made, but one possibility is that the 12 DCB's in a subrack may share eight 12-channel

Molex parallel optical links with a bandwidth in excess of 2.5 Gbps per optical fiber. Each DCB would transmit data on 8 fibers with an aggregate bandwidth in excess of 20 Gbps.

10.2.5 Initialization, Control, and Monitoring

Slow control and monitoring of the FPIX2 pixel readout chips is accomplished through the FPIX2 programming interface, which operates independently of the data output interface. All FPIX2's on a single HDI share a synchronous serial programming link. The BCO clock is used to clock this link. Individual FPIX2 chips are identified by a 5-bit chip id, which is set by internal wire bonds. Commands to write to registers can be sent from the DCB to individual chips or broadcast to all the chips that share a programming link. Read commands must be sent to a single chip.

At the beginning of data taking, a variety of internal FPIX2 registers must be set. These include registers that control the operation of each chip as a whole, such as the discrimination thresholds and internal bias voltages and currents, and two registers that control individual pixel cells (pixel kill and test charge inject). All registers can be reset to default settings with a single command.

All control registers (except kill and inject) can be read non-destructively. During data taking, registers are periodically read back and their contents checked. If a bit error is detected (the Single Event Upset rate has been measured to be very low, but is non-zero), the register is reset. Most errors can be corrected without halting data acquisition.

Pulser calibration data will be taken during beam gaps and other times when no collisions are occurring. These data will be used to verify operational parameters such as discrimination levels and lists of dead and hot pixels.

10.3 RICH Detector

10.3.1 Overview

The RICH detector includes two different subsystems: the gas RICH, which has two viable solutions for the photon detector system, either hybrid photon detectors (HPD's) or multi-anode photomultiplier tubes (MAPMT's), and the liquid radiator RICH, which uses 3" photomultiplier tubes as photon detectors. These subsystems have very different analog signal properties, but their readout has the same conceptual design. Photodetector elements are connected to the readout electronics hosted in the front end hybrids (FE-HYB) described in more detail in the RICH detector section of the TDR. Several hybrids are connected with a front end multiplexer board (FE-MUX) that provides the interface between these front end hybrids and the associated data combiner board (DCB).

The HPD readout system consists of 153,872 readout channels grouped into 944 front end hybrids: each hybrid processes the signals from the 163 HPD pixels. More details on the technology used to implement the front end hybrids and the custom made ASICs (VA_BTeV) is available in the RICH TDR section. The grouping of front end hybrids into

FE-MUX boards parallels the mechanical grouping of HPD's. Signals from up to 6 HPD hybrids are combined in a single multiplexer board that provides local buffering of the HPD hits, time stamps and bidirectional data transfer between the DCB's and the FE-HYB's. The design of the FE-MUX boards is such that they can be used with the cables and DCB design adopted for other BTeV subdetectors. The alternative approach, based on MAPMT's, includes a similar number of readout channels (144,256). Each MAPMT includes 16 pixels. 8 MAPMT share the same front end hybrid, similar in conceptual design to the HPD front end hybrid. Also in this case front end multiplexer boards provide the interface between the RICH front end hybrids and the DCB's. Finally, the liquid radiator system includes 5048 PMT's. We plan to read them out with the same front end ASIC's developed for the MAPMT's.

10.3.2 Component Quantities

Table 10.2 summarizes the number of front end hybrids, multiplexer boards, DCB's, and serial link data cables used in the RICH readout system. A block diagram of the system is given in Figure 10.3.

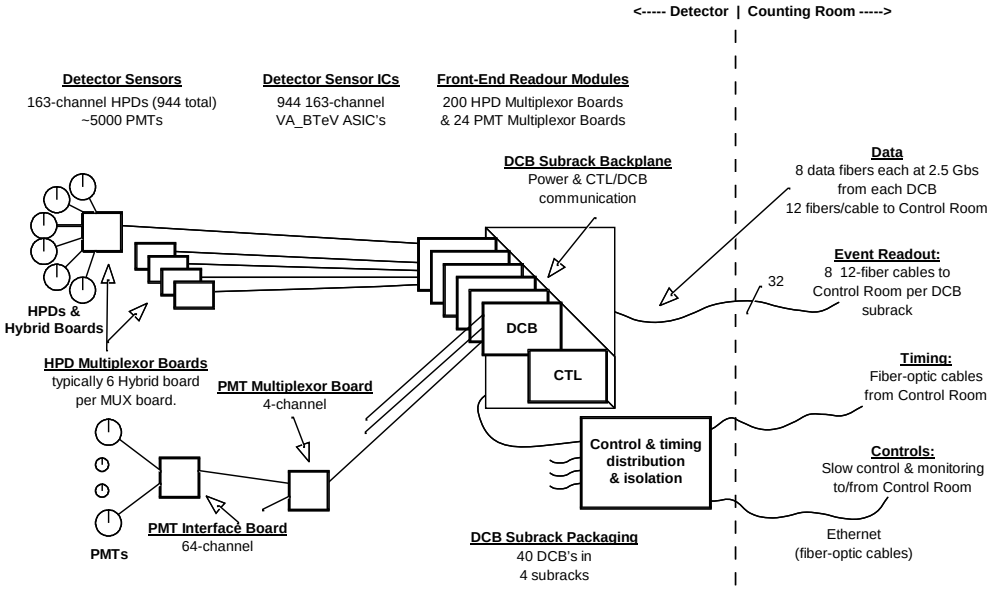

Figure 10.3: Block diagram of BTeV RICH Detector front-end electronics (shows the HPD option for the gas RICH).

Gas RICH	FE hybrids	FE MUX	Data link cables	DCB's
HPD Option	944	200	517	30
MAPMT Option	1196	332	682	36
Liquid RICH	FE hybrids	FE MUX	Data link cables	DCB's
PMT	80	24	58	4

Table 10.2: FE Modules, DCB's, and I/O cables required for the RICH detector subsystems (assuming a 16-bit data word is used).

10.3.3 Data structure

The output data from each front end hybrid is a bit stream containing one bit of information for every detection element for a given beam crossing. The conceptual design will be discussed with reference to the HPD system. In this case, each front end hybrid generates a 163-bit data word for each crossing. In order to use standard 50 pin connectors and cables and minimize interconnection cost, this information is shipped in 5 bursts, using a 151 MHz clock, phase locked with the BCO clock, that is provided by the DCB's to the FE-MUX boards.

The multiplexer board receives this information and generates the output data to be shipped to the DCB's. A FPGA in the front end multiplexer board takes the information from the 6 front end hybrids and stores the hit channel addresses in a FIFO memory. For each bunch crossing, the first step is to save in the FIFO the event data block (sparsified addresses). The format for the data block has not yet been finalized. If the data is stored in 18-bit words, then the only header information required is a BCO time stamp. If 16-bit words are used, header records will also be required to indicate which section of the detector the following data belongs to. In either case, the header records are followed by a list of addresses identifying the channels recording a hit during the specified BCO.

Data transfer from the FIFO to the DCB uses a 500 Mbps serial data output link also used by many of the other BTeV detector components. A slow control and timing link is included in the same cable bundle. The number of data links that are needed for each front end multiplexer board is strongly affected by the location of the photosensitive detector. The occupancy is quite different at the various HPD locations, as discussed in more detail below. The number of DCB's and data links required is also slightly dependent on the data format chosen. Our present design incorporates a variable number of serial lines depending upon the location of the front end multiplexer board. This system is suitable for any beam configuration envisaged so far, including effects of electronics noise and 20% excess capacity to account for unforeseen effects. The DCB's collect information from several FE-MUX boards and send the information, suitably grouped and formatted, to the data acquisition buffers.

10.3.4 Occupancy and Data Rate Studies

An extensive discussion of the occupancy studies that we performed, both in the case of 132 ns and 396 ns beam crossing period and with different assumptions for the number of interactions, has already been presented in the RICH detector description. Here we summarize the work done to determine that the readout architecture is suitable even for the hottest region in the gas RICH detector. To this purpose we have combined our physics simulation with a model of the readout system.

We have used BTeV Geant to simulate elastic and inelastic collisions. For the 132 ns scenario, we have generated events with a number of interaction per crossing following a Poisson distribution with mean 2. For the 396 ns scenario we have followed a similar approach, with a mean value of 6 interaction per crossing. We have added to the physics hits an extra 1% of noise hits. For each event, data blocks are stored in a local FIFO and then shipped to the DCB's with multiple serial lines as described before.

At 2 interactions per crossing and 132 ns crossing period, the hottest group of HPD's registers an average of 66.2 hits per crossing, of which 57 are induced by Cherenkov photons and 9.2 are noise hits. Assuming 16-bit words, this corresponds to an average data rate from the hottest multiplexer card of 8 Gbps. This multiplexer will be serviced by 20 output links with an aggregate bandwidth of 10 Gbps. If the Tevatron operates with 396 ns between crossings, then the hottest group of HPD's will register three times as many Cherenkov photons (171) but essentially the same number of noise hits (8). Thus, a readout solution which works for 132 ns operation will be more than adequate for 396 ns operation.

These simulations show that the overwhelming majority of the data produced by the RICH is produced by the HPD's. With 132 ns between crossings, and an average of two interactions per crossing, the total number of HPD hits due to light produced by tracks is 1040 per crossing. The estimated electronic noise (1% of all channels) adds another 1528 hits, so the total number of hits per crossing is 2568. Assuming a 16-bit word data format, and ignoring the small contribution to the total amount of data due to header words, the aggregate data rate is estimated to be:

$$\frac{2568 hits}{132 ns} \times \frac{16 bits}{hit} = 0.31 Tbps. \tag{10.2}$$

With 396 ns between crossings and an average of six interactions per crossing, the number of hits produced by tracks triples, but the number of noise hits remains constant, so the total number of hits per crossing is only 4943. This implies a slightly lower total data rate of:

$$\frac{4943 hits}{396 ns} \times \frac{16 bits}{hit} = 0.20 Tbps. \tag{10.3}$$

10.3.5 Initialization, Control and Monitoring

Prior to data taking, an initialization sequence needs to be run to set the operation mode of the front end ASICs (e.g. calibration or normal run), the list of channels that are enabled and fine-tune the threshold of individual channels. Moreover there are some analog voltages

that need to be set. The initialization sequence requires a clock and a serial line that shifts the information in the relevant registers. Commands to write the registers will be sent from the DCB's. A serial read-back of the initialization sequence can be performed for monitoring purposes.

During data taking several parameters will be monitored. In particular, the temperature on all the front end hybrids and key voltages and currents will be monitored through the slow control system. Additional quantities monitored include the expansion volume, the gas and liquid recirculation system and the cooling system. The RICH slow control will be implemented in the framework of the BTeV control and monitoring system.

Periodic calibrations will be performed during times when no collisions are occurring using a pulser to inject a controlled amount of charge on a calibration capacitor located on the front end board. In addition LEDs interspersed in the detector volume allow calibration of the overall photodetector system.

10.4 Electromagnetic Calorimeter

10.4.1 Overview

The Electromagnetic Calorimeter is comprised of approximately 10,100 lead tungstate crystals. The crystals are 28mm × 28mm × 220mm and are arranged in a fashion that results in a circular array with a radius of 1.6 meters. Bonded to the back of each crystal is a photomultiplier tube that detects the light generated in the crystal and converts the light to an electrical signal.

Power is provided to the photomultiplier tubes by high-voltage power supplies located outside of the collision hall. Connections between the PMTs and the power supply cables are accomplished with circuitry on PMT bases. Multiple voltages are used for each PMT and groups of approximately 100 PMTs are ganged together, sharing common power supply channels.

Analog signals from the photomultiplier tubes are sent via copper cable to custom electronics located in 20 electronics subracks positioned in racks near the calorimeter. Each Analog-to-Digital Converter (ADC) Card has 32 channels of electronics. This set of electronics digitizes the analog signals from the PMTs and performs zero suppression on the data. The analog-to-digital conversion is accomplished with a full-custom application-specific integrated circuit (ASIC) that is a new version in a series known as the QIE (charge Integrating and Encoding) chip developed at Fermilab. There are 16 ADC Cards per subrack. Power distribution and board-to-board communication within the subrack is accomplished with a custom backplane. The design and construction of these racks, or the cable connection methodology between the racks and the detector, allows the calorimeter to move 16 inches in the z-direction.

Data generated by the ADC Cards is sent to Data Combiner Boards which concentrate and pass the data on to the remaining portions of the data acquisition system. The Data Combiner Board also provides synchronization signals to the ADC Card and provides a path

for signals implementing slow control and monitoring functions. Data is sent from the ADC cards to the DCB's using the same 500 Mpbs serial links used in a number of other BTeV subsystems. A 32-channel ADC card can be configured to use as many as 16 500 Mbps links or as few as one link. 16 links provides enough bandwidth so that no zero suppression is required, even with a time between crossings of 132 ns. ADC cards servicing crystals very close to the beam will be configured with 16 data output links. Those servicing crystals further from the beam will be configured to use fewer links. Most of the ADC cards will require only one 500 Mbps output link.

The Data Combiner Boards used for EMCal readout will be very similar, if not identical, to those used by many of the other BTeV detector subsystems. Each DCB will accept inputs from up to 48 500 Mbps data links on 24 separate cables. The number of EMCal DCB's will be determined by cabling and packaging convenience and by the number of data link cables required. Our present design calls for 24 DCB's.

10.4.2 Component Quantities and Locations

PMT's	10100
32 channel Transition Cards	316
32 channel ADC Cards	316
QIE9 ASIC's	10100
ADC Subracks	20
DCB's	24
DCB Subracks	2

Table 10.3: Component Count

A block diagram of the Electromagnetic Calorimeter front end electronics is given in Figure 10.4. Table 10.3 summarizes the number of components used in this system.

10.4.3 Data Structure

The data word generated by a QIE9 will include an 8-bit mantissa, a 3-bit range (exponent), and a 2-bit capacitor id. This output will be compared with a digital threshold on the ADC card. Only channels that are above a programmable threshold will be read out. For each crossing, the ADC card will create a header word containing a beam crossing number (time stamp) and possibly a count of the number of values being read out. Hit data will consist of a channel number and a 13-bit QIE value. If 16-bit data words are used, then 2 words of data will be required for each hit.

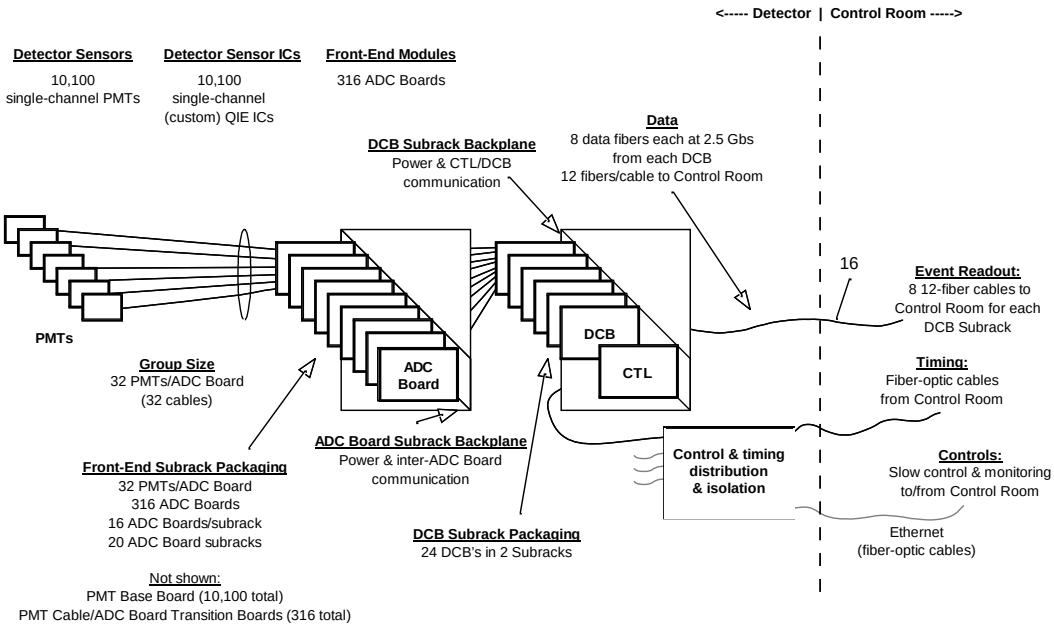

Figure 10.4: Block diagram of BTeV Electromagnetic Calorimeter front-end electronics.

10.4.4 Occupancy and Data Rate Estimate

BTeV GEANT simulations indicate that at the BTeV design luminosity of 2×10^{32} cm$^{-2}sec^{-1}$ approximately 1000 hits will need to be read out every 132 ns, or approximately 3000 hits every 396 ns, depending on the operating mode of the Tevatron. In either case, the data rate will be approximately eqaul to:

$$\frac{1000hits}{132ns} \times \frac{32bits}{hit} = 0.24Tbps. \tag{10.4}$$

If the data load is shared equally by the DCB's, then the rate out of each of the 24 DCB's will be approximately 10 Gbps. This represents approximately one half of the output bandwidth planned for the "standard" DCB's.

10.4.5 Initialization, Control, and Monitoring

Slow control and monitoring of the ADC cards will be accomplished using the control path associated with the ADC serial data output link. This link is described in more detail in descriptions of the Straw Detector readout. Before data taking, digital thresholds will be downloaded to the ADC cards. A synchronous reset signal will insure that all of the QIE9 chips start in a known state (it is necessary to know which pipeline capacitor is associated with which beam crossing).

Calibration of the QIE chips and associated electronics will be accomplished (during periods when no collisions can occur) with a DC current source that delivers a known amount of charge to the QIE inputs.

10.5 Muon Detector

10.5.1 Overview

The basic building block in the construction of a detector station is a "plank" of 3/8" diameter stainless steel proportional tubes. There are 32 tubes in each plank, arranged in two rows of 16 offset by half a tube diameter. These are held together with aluminum ribs and by the brass gas manifolds which are glued to the end of each plank. Each plank is a sturdy, self-supporting building block which acts as an excellent Faraday cage. All the tubes in the plank are terminated on one end and read out on the other. There are a total of 1152 planks in the muon detector. The data from each plank is sparsified at the detector by an FPGA (or equivalent) and sent to DCB's using LVDS over copper serial links. Each serial link consists of a single cable with an RJ45 connector, capable of supporting 2 additional links, one of which will be used for slow control. These links are identical to those which will be used to carry data to the DCB's for the Forward Straw detector, the Electromagnetic Calorimeter, and the RICH detector.

To minimize occupancy at small radii, planks of increasing length are arranged into pie shaped octants. To minimize pattern recognition confusion, three arrangements of planks (r, u, or v) are used. The r views are radial. The u and v views are rotated ± 22.5 degrees with respect to the radial views and measure the azimuthal angle, ϕ. A collection of 8 octants of like arrangement is called a view, and a collection of 4 views is called a station. In order to provide redundancy in the most important view in terms of pattern recognition for the trigger and momentum measurement, the r view is repeated in each station. The whole muon detector is three stations located at the end of the BTeV experiment, interspersed between and after magnetized iron toroids and shielding. A block diagram of the front end electronics is shown in Fig. 10.5.

10.5.2 Component Quantities and Locations

ASDQ chips will perform the analog to digital conversion of the proportional tube signals. For our purposes, each ASDQ chip consists of 8 channels of amplifier, shaper and discriminator with a common threshold. In Table 10.4, we summarize the number of planks, ASDQ readout chips, FPGA's, DCB's, DCB subracks, and data links used in the muon detector.

10.5.3 Data Structure

The data from the muon front end system will have 2 formats depending on the number of hits in the plank. A data header consists of a 12 bit plank ID consisting of 4 bits(plank)

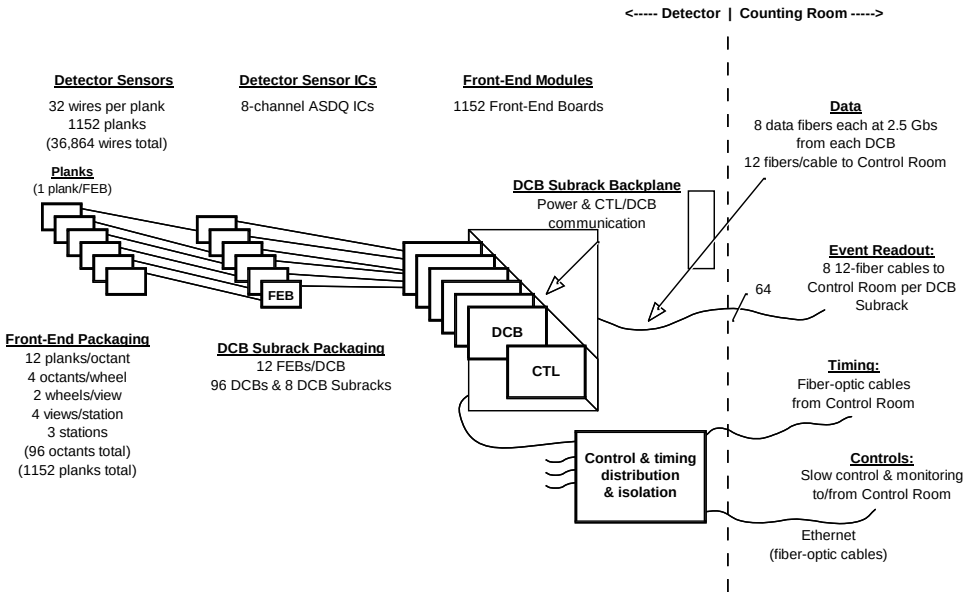

Figure 10.5: Block diagram of BTeV Muon Detector front-end electronics.

	Number per "plank"	Number per "octant"	Number per detector
Tubes	32	384	36,864
ASDQ	4	48	4608
FPGA	1	12	1152
DCB's	1/12	1	96
DCB subracks	1/192	1/16	6
500 Mbps data links to DCB's	1	12	1152

Table 10.4: Component Count

+3 bits(octant) +2 bits(view)+2 bits (station), an 8 bit beam crossing number, a 1 bit data type and either a 3 bit word count followed by up to 5, 5 bit words, or a 32 bit hitmap. Since the data coming from the muon system is easily time ordered, the DCB's can optimize sparsification over the entire octant before sending the data to the Level-1 Buffers and the muon trigger.

	Station 1	Station 2	Station 3	Total
Average number of hits per crossing	42	8	9	54
Average occupancy	0.34%	0.06%	0.07%	0.15%
Maximum channel occupancy	2.5%	0.24%	0.52%	
Maximum plank occupancy	1.6%	0.17%	0.31%	

Table 10.5: Muon detector occupancies obtained from BTeVGeant simulations with an average of 2 minimum bias interact ions per crossing and a crossing rate of 7.6 MHz (132 ns bunch spacing). Average occupancy is the occupancy of the detector in a single crossing. Maximum channel occupancy is the maximum hit rate for the innermost channel. Maximum plank occupancy is the average per channel hit rate of the innermost plank.

10.5.4 Occupancy and Data Rate Estimate

At a nominal luminosity of $2 \times 10^{32}\text{cm}^{-2}\text{s}^{-1}$ and a bunch spacing of 132 ns, we expect 2 minimum bias interactions/crossing which are simulated using a Poisson distribution with mean of two. In Table 10.5 we summarize the detector occupancies obtained from BTeVGeant under this scenario. These occupancies and rates are fairly low by modern detector standards, and will remain low even if the interactions/crossing is increased.

Our estimate for the highest data average data rate coming from a single plank is estimated to be:

$$0.016(Average\ Occupancy) \times 32(chnls) \times 28(bits) \times 15.2 \times 10^6(\frac{interactions}{sec}) = 0.22\ Gbps$$

Our estimate for the highest average data rate into into a single muon DCB, under the assumption that each hit is unique in a plank, is expected to be:

$$0.0034(Average\ Occupancy) \times 384(channels) \times 28(bits) \times 15.2 \times 10^6(\frac{interactions}{sec}) = 0.56\ Gbps$$

Our estimate for the highest average data rate coming from the muon system, under the assumption that each hit is unique in a plank, is expected to be:

$$54(Average\ \ total\ hits) \times 28(bits) \times 15.2 \times 10^6(\frac{interactions}{sec}) = 23\ Gbps$$

These estimates do not take into account noise in the detector. From our prototype tests, the noise/plank is small, on the order of 10 hz/tube. This adds an additional:

$$1152(channels) \times 28(bits) \times 10/sec = 0.32\ Mbps$$

to the data rate coming from the detector.

10.5.5 Installation, Control and Monitoring

Slow control, monitoring and data output from the muon front end boards is accomplished through the FPGA on the card. Slow serial links will be used for monitoring and control, and fast links will be used for the data output.

The data output will be gated at a nominal width of 120 ns. In preliminary tests it was determined that the arrival time of a data pulse could be localized to within 5 ns inside this gate, adding the possibility that a TDC function could be included in the FPGA programming. In order to form an output data word identifier, we assign an ID for each plank. Each front end card has a 16 bit chip ID which is linked to the bar code placed on the card during construction. All the test data from the ASDQ's on the card, the card itself, the plank, and the individual tubes in a plank will be linked in a database and tracked during construction. Placement of a plank in an octant will also be stored during construction and checked against the internal ID during the octant test. The location will then be used to assign the correct 12 bit sequence, stored in non-volatile memory, for the data word attached to the hits during readout. (The 16 bit card ID can be used as well, and the 12 bit data word ID can be assigned in the DCB in the event of a malfunction.)

A programmable default configuration will be set for the card which can be invoked with a reset to the card. The configuration of the card will be periodically checked via the slow control and a reset issued if needed. In addition to the monitoring, the setting of the threshold DAC's and various other control lines to the ASDQ, the slow control is used to invoke a test pulse common to each ASDQ and synchronized to the beam crossing clock. The beam crossing clock is delivered to the front end via the other serial line reserved for slow data.

10.6 Forward Straw Detector

10.6.1 Overview

The basic building block of the forward straw detector is the straw module. A straw module consists of 48 straws. Straw modules are combined to make up a view, and 3 views (wires vertical and tilted at plus and minus 30 degrees from vertical) make up a station. Seven stations are spread out longitudinally along the beam from station 1 (nearest the interaction point) to station 7 (furthest).

High voltage distribution, hit detection and time-to-digital conversion are performed by an "electronics package." The sense wires in high occupancy straw modules have a glass bead in the middle, effectively dividing the wire in two. These straw modules are serviced by two electronics packages; lower occupancy modules are serviced by one.

ASDQ ASIC's amplify and discriminate the straw anode wire signals. A Fermilab-designed 24-channel TDC ASIC measures the drift times. Each electronics package communicates with the Data Combiner Board (DCB) using 4 differentially driven, twisted pair (copper) serial links, two used for communications from the DCB's and two used for data

from the electronics package. A 132-ns Reference Clock is used to generate all precision clocks used within the electronics package, such as the beam crossing clock. A 151.1 Mbps "Timing and Control" link defines beam crossing markers and initiates CSR read/write functions. Two "event data" links, operating at 636 Mbps (with 8b/10b encoding and a bandwidth of 509 Mbps) move data from the electronics package to the Data Combiner Board. One data link is devoted to each of the two TDC ASIC's in the electronics package.

Commercially available receiver-equalization chips are used to compensate for the frequency dependent characteristic of the copper cables and restore the eye diagram of data on the event data links. Commercially available driver-end pre-emphasis chips are used to partially correct the same frequency dependent characteristic for the Reference Clock and Timing/Control links. Note that these components are located away from the straw modules, on the DCB.

10.6.2 Component Quantities and Locations

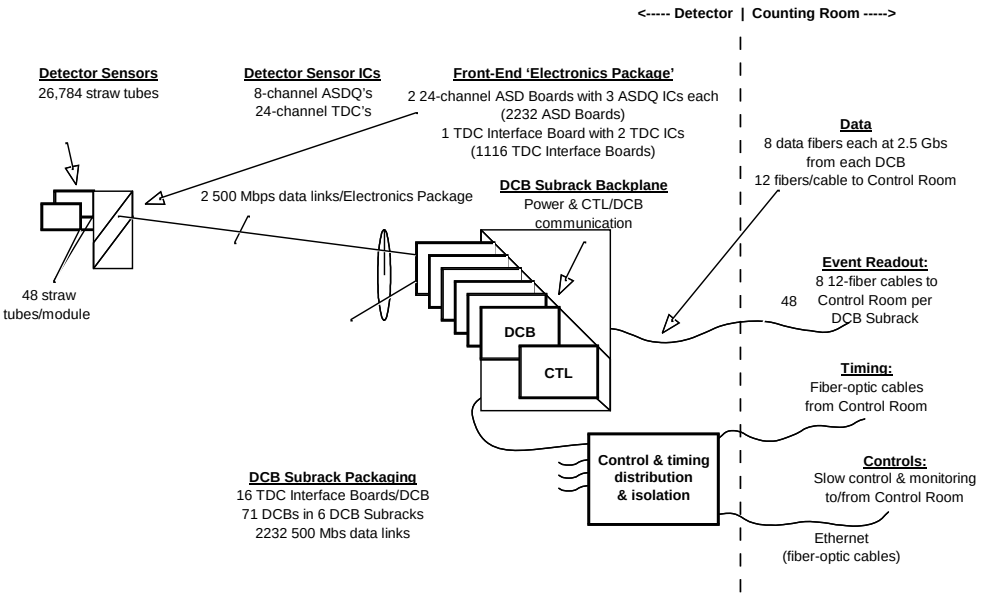

Figure 10.6: Block diagram of BTeV Forward Straw Detector front end electronics.

Table 10.6 summarizes the number of straw modules, electronics packages, ASDQ's, TDC's, and DCB's used in the Forward Straw Detector. A block diagram of the system is given in Figure 10.6.

Station (3 views)	Modules	Straws	Electronics Packages	TDC Cards	TDC Chips	ASDQ Cards	ASDQ Chips	DCB's
1	24	1152	48	48	96	96	288	3
2	36	1728	72	72	144	144	432	5
3	54	2592	108	108	216	216	648	7
4	78	3744	156	156	312	312	936	10
5	90	4320	180	180	360	360	1080	11
6	102	4896	204	204	408	408	1224	13
7	174	8352	348	348	696	696	2088	22
Total	558	26784	1116	1116	2232	2232	6696	71

Table 10.6: Component Count

10.6.3 Data Structure

The format of the data word generated by the TDC ASIC has not yet been specified. Each TDC ASIC will generate a "header" word every 132 ns which will ensure that data is assigned to the correct crossing number. In the data rate estimates given below, we assume that this header word is 16-bits long. Drift times will be encoded using 5-7 bits. Another 5 bits are required to specify which of 24 wires is hit. If each hit is packed into a 16 bit word, a few more bits may be used in the wire number to minimize the amount of data reformatting which must be done in the DCB's.

10.6.4 Occupancy and Data Rate Estimate

Simulations have been done on the expected occupancy of all modules in the detector. The required output bandwidth (assuming 16-bit data words) for the worst-case module in each station, and for the station as a whole, is shown in Table 10.7. The occupancies listed were calculated assuming a luminosity of 2×10^{32} cm$^{-2}sec^{-1}$ and a time between crossings of 132 ns. If the time between crossings is 396 ns, the worst case occupancies will be slightly less than three times higher. The data rate out of the front end electronics will be essentially unchanged.

The aggregate data rate output from the straw system DCB's to the Level-1 buffers can be estimated as follows. The average number of hits per crossing, given two interactions per crossing, is approximately 1050. The data stream will contain many fewer header records that the data stream input to the DCB's, since packets will be sent only for crossings containing hit data, and many front end data streams will be merged. Ignoring header records (and using one 16-bit word per hit)the data rate out of the straw system DCB's is estimated to be:

$$\frac{1050 hits}{132 ns} \times \frac{16 bits}{hit} = 127 Gbps. \tag{10.5}$$

Station	Maximum Occupancy	Maximum Bit Rate	Bit Rate for Station
1	4.1%	240 Mbps	21.2 Gbps
2	4.9%	264 Mbps	30.2 Gbps
3	5.8%	290 Mbps	42.0 Gbps
4	6.2%	302 Mbps	54.7 Gbps
5	6.6%	313 Mbps	62.3 Gbps
6	6.6%	313 Mbps	68.7 Gbps
7	8.0%	354 Mbps	119.5 Gbps
Total			400 Gbps

Table 10.7: Worst case occupancy and data rates

10.6.5 Initialization, Control, and Monitoring

The TDC ASIC's will include digital to analog converters. Some of these DAC's will be used to provide discriminator threshold voltages for the ASDQ chips. The threshold values will be reloaded at the beginning of each store. Two test modes are envisioned for the data link between the TDC's and the DCB's; a loopback test in which data is sent from the DCB's to the TDC's and back to the DCB's, and a simpler test in which the TDC's generate a fixed pattern which can be verified by the DCB's. The TDC's will also include a self-test feature in which pulses will be input to each channel at known times. Monitoring of configuration data, and link and TDC tests will occur between stores and perhaps in the beam abort gaps.

10.7 Forward Silicon Detector

10.7.1 Overview

The Forward Silicon Detector includes seven identical stations, each with three planes, all of identical construction. Each plane comprises four ladders, and each ladder in turn is made up of four sensors. Strips on the sensors closer to the center of the ladder are wire bonded to the outside sensors and connected to FSSR readout chips located at both ends of the ladders. The FSSR readout chips share an architecture with the FPIX2 pixel readout chips. The FSSR's communicate with DCB's using LVDS over copper serial links. All of the FSSR chips on one end of a ladder share one slow control and monitoring link, as well as common digital and analog voltage and ground. Hit data is output from the FSSR chips on 140 Mbps point-to-point links. Chips on ladders close to the beam are configured to use four data output links. Those on ladders further from the beam are configured to use two data output links.

Signals are carried on lightweight flexible cables between the ends of the ladders and "splitter boards" located outside the active area of the straw chambers. Signals are carried between the splitter boards and the DCB's on conventional high-density cables ($\sim$ 5 m long). The DCB's are located in racks mounted as close as practical to the silicon strip stations. Each DCB is connected to 12 half ladders, so there are two DCB's per station.

10.7.2 Component Quantities and Locations

Table 10.8 summarizes the number of ladders, sensor wafers, readout chips, DCB's, DCB subracks, and data links used in the forward silicon detector.

	Number per plane	Number in complete detector
4-sensor ladders	4	84
sensor wafers	16	336
FSSR IC's	48	1008
DCB's	see text	14
DCB subracks	see text	2
140 Mbps data links to DCB's	144	3024
2.5 Gbps data links from DCB's	NA	168

Table 10.8: Component Count. Each of the seven FSIL stations contains three planes. Each ladder has 700 strips read out per end and is instrumented with six FSSR readout chips on each end.

10.7.3 Data Structure

The silicon strip hit format is exactly the same as the pixel hit formal. This choice of format was made to simplify the design of the FSSR readout chip, even though at least 7 of the 23 bits are not needed to encode silicon strip data. Like the pixel DCB's the silicon DCB's add a 7 bit chip number to each data word. Also, the data from the FSSR's, like data from the FPIX2's, is not strictly time-ordered. As in the pixel system, time order will be restored by the modules in the control room that receive the silicon data. It is not yet determined whether or not the silicon DCB's need to extend the beam crossing number before transmitting data to the counting room, or whether or not the silicon DCB's will reformat the data recieved from the FSSR's to eliminate unused bits.

10.7.4 Occupancy and Data Rate Estimate

In minimum bias interactions generated by Pythia and simulated by BTeV GEANT, an average of 97 hits are generated in the Forward Silicon Detector per interaction. Assuming

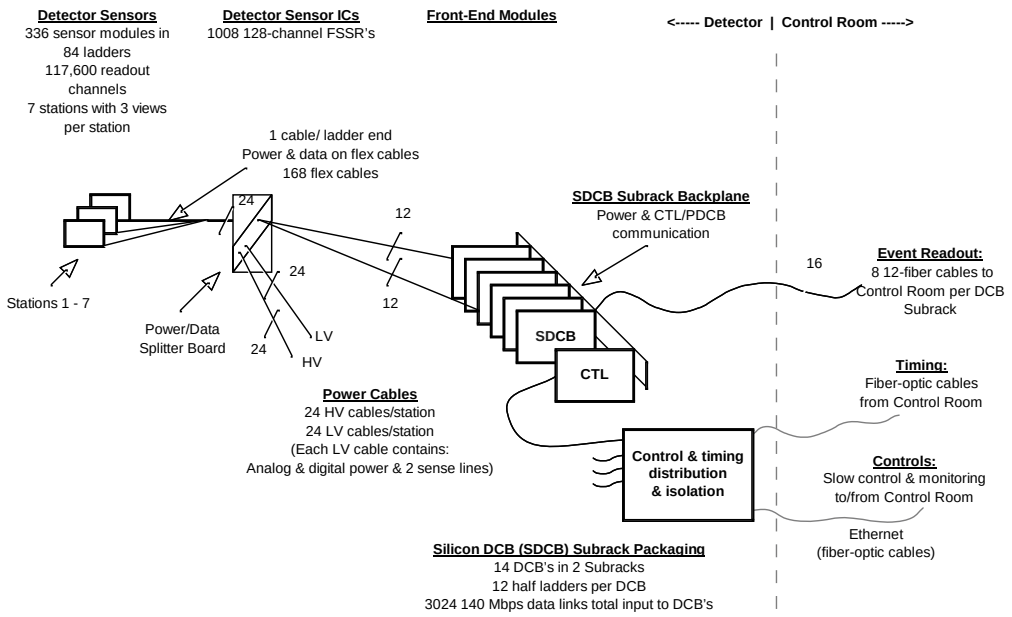

Figure 10.7: Block diagram of BTeV Forward Silicon Detector front-end electronics.

that 70% of these result in one silicon strip hit, and 30% result in two-strip clusters, the number of silicon strips hit per interaction is approximately 130. Ignoring noise hits, the data rate into the FSIL DCB's at the BTeV design luminosity of 2×10^{32} cm$^{-2}sec^{-1}$ is estimated to be:

$$130\frac{hits}{interaction} \times 24(\frac{bits}{hit}) \times 15.2 \times 10^6(\frac{interactions}{sec}) = 47Gbps. \qquad (10.6)$$

This is to be compared with the aggregate output bandwidth of the FSSR readout chips of approximately 0.4 Tbps. The silicon strip DCB's, like the pixel DCB's, may add information to each hit before transmitting it to the counting house over a commercial optical link. Nonetheless, an aggregate bandwidth of slightly less than 280 Gbps (which would be provided using 8 2.5 Gbps links per DCB) between the DCB's and the control room will insure that these links do not become data bottlenecks.

10.7.5 Initialization, Control, and Monitoring

The slow control and monitoring interface of the FSSR silicon strip readout chips is identical to that of the FPIX2 pixel readout chips. All FSSR's on one end of a ladder share a synchronous serial programming link. Control and monitoring of the FSSR's will be very similar to control and monitoring of the FPIX2's.

10.8 Data Rate Summary

The data rate into the DCB's and the number of DCB's used for each of the detector subsystems is summarized in Table 10.9.

Detector Subsystem	Data Rate into DCB's	Number of DCB's
Pixel Detector	300 Gbps	120
RICH Detector	200-310 Gbps	40
EM Calorimeter	240 Gbps	24
Muon Detector	23 Gbps	96
Forward Straw Detector	400 Gbps	71
Forward Silicon Detector	47 Gbps	14
Total	1.2-1.3 Tbps	364

Table 10.9: Data Rate into the Data Combiner Boards